Attachment A

Application Cover Sheet

Name of Respondent: Element Solutions LLC			
Mailing Address 14 Ridge Square NW	City/Town Washington, DC	State DC	Zip Code 20016
Telephone 240-463-8176	Fax N/A	Web Address: e14s.com	
Primary Contact for Clarification Nicolas Ojeda		Primary Contact E-mail Address Nico@e14s.com	
Authorized Signatory Nathaniel Barbato		Authorized Signatory E-mail Address nathaniel.barbato@e14s.com	
Legal Status/Jurisdiction(e.g., a Massachusetts Corporation, LLC, LLP, etc.) LLC, MD		Respondents W-9 Respondents UEI No. H71DAHWDKN64	

Attachment B

Massachusetts Technology Collaborative Authorized Respondent's Signature and Acceptance Form

The undersigned is a duly authorized representative of the Respondent listed below. The Respondent has read and understands the RFP requirements. The Respondent acknowledges that all of the terms and conditions of the RFP are mandatory, and that Respondent's response is compliant with such requirements.

The Respondent understands that, if selected by the Mass Tech Collaborative, the Respondent and the Mass Tech Collaborative will execute an Agreement specifying the mutual requirements of participation.

Respondent agrees that the entire bid response will remain valid for sixty (60) days from receipt by the Mass Tech Collaborative.

I certify that Respondent is in compliance with all corporate filing requirements and State tax laws.

I certify that the statements made in this response to the RFP, including all attachments and exhibits, are true and correct to the best of my knowledge.

I acknowledge and understand the procedures for handling confidential materials submitted to the Mass Tech Collaborative as set forth in Section 3.1(d); I agree to be bound by those procedures and agree that the Mass Tech Collaborative shall not be liable under any circumstances for the disclosure of any materials submitted to the Mass Tech Collaborative pursuant to this RFP or upon the Respondent's selection.

Respondent: Nathaniel Barbato

By: 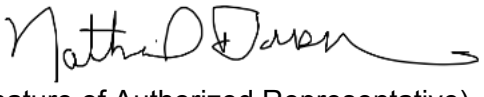
(Signature of Authorized Representative)

Name: Nathaniel Barbato
Title: Director
Date: 10/06/2025

element

Application Narrative

AI Hub Data Commons Collaborative Consulting Services

RFP No. 2026-JAII-01
Date October 07, 2025

Submitted by:

Amit Thakkar, CEO
Amit@e14s.com

Nathaniel Barbato, Director
Nathaniel.Barbato@e14s.com

Company data:

14 NW Square, Suite 300
Washington, DC 20016

SAM UEID: H71DAHWDKN64

This proposal contains **proprietary and confidential** information of Element Solutions which is disclosed to the recipient for the sole purpose of evaluating our Technical Proposal Narrative Massachusetts AI Hub Data Commons Collaborative Proposal and may not be used, reproduced, or disclosed to any third party without prior written consent.

1. Executive Summary

The Massachusetts Technology Collaborative (MassTech) is advancing a bold vision to establish the Commonwealth as a national leader in ethical and responsible AI. At the heart of this vision is the Data Commons Collaborative (DCC), a secure, FAIR-compliant (Findable, Accessible, Interoperable, Reusable), metadata-driven environment designed to power innovation across healthcare, climate, transportation, robotics, and beyond. Element Solutions and JDSAT (Team Element) are proud to submit a solution that aligns to MassTech's requirements while offering a proven, phased roadmap for delivery and adoption.

element Element delivers secure, standards-based digital platforms that transform how governments manage and share data. We break down complex challenges into modern, user-focused solutions that drive efficiency, maximize taxpayer value, and deliver measurable outcomes. With deep expertise in data commons, metadata-driven architectures, and governance, we design scalable, cloud-agnostic platforms that integrate FAIR principles, synthetic data, and AI fairness frameworks. While our foundation is in healthcare, our innovation extends across government, where we streamline operations, strengthen compliance, and empower agencies to harness data responsibly for public impact.

JDSAT JDSAT delivers advanced analytics and AI solutions that empower organizations to transform complex data challenges into actionable insights and decision advantage. Our expertise spans operational decision analysis, strategic analytic advice, and digital transformation, with specialized capabilities in synthetic data generation and AI fairness frameworks. Drawing on our experience supporting federal agencies with data governance and cloud-native infrastructure, JDSAT offers the Massachusetts AI Hub a trusted partner with proven methodologies for implementing secure, scalable data platforms that advance innovation while maintaining ethical standards. Our team of data scientists, engineers, and domain experts stands ready to deliver the technical architecture and governance frameworks necessary for a successful Data Commons Collaborative.

Our approach combines Element's record of delivering complex federal data platforms with JDSAT's expertise in synthetic data generation, fairness frameworks, and AI integration. The proposed solution unifies CKAN and Dataverse to provide a research-grade platform that enforces interoperability, provenance, and FAIR compliance. From launch, the platform will be WCAG 2.1–accessible, multilingual, and cloud-agnostic. Security and governance are embedded by design, leveraging Zero Trust architecture, StateRAMP, HIPAA, NIST 800-53, and Massachusetts privacy laws to ensure compliance and public trust.

Team Element designed a phased deployment plan that delivers measurable outcomes early while scaling responsibly over time. Each phase provides stakeholders with clear, incremental benefits: early access to curated datasets, safe synthetic data for experimentation, bias monitoring for equitable outcomes, and transparent governance for trust and accountability. Our Commons Credits Ledger adds a

unique sustainability dimension by incentivizing participation and rewarding carbon-efficient AI practices.

The following graphic illustrates Team Element's phased deployment approach.

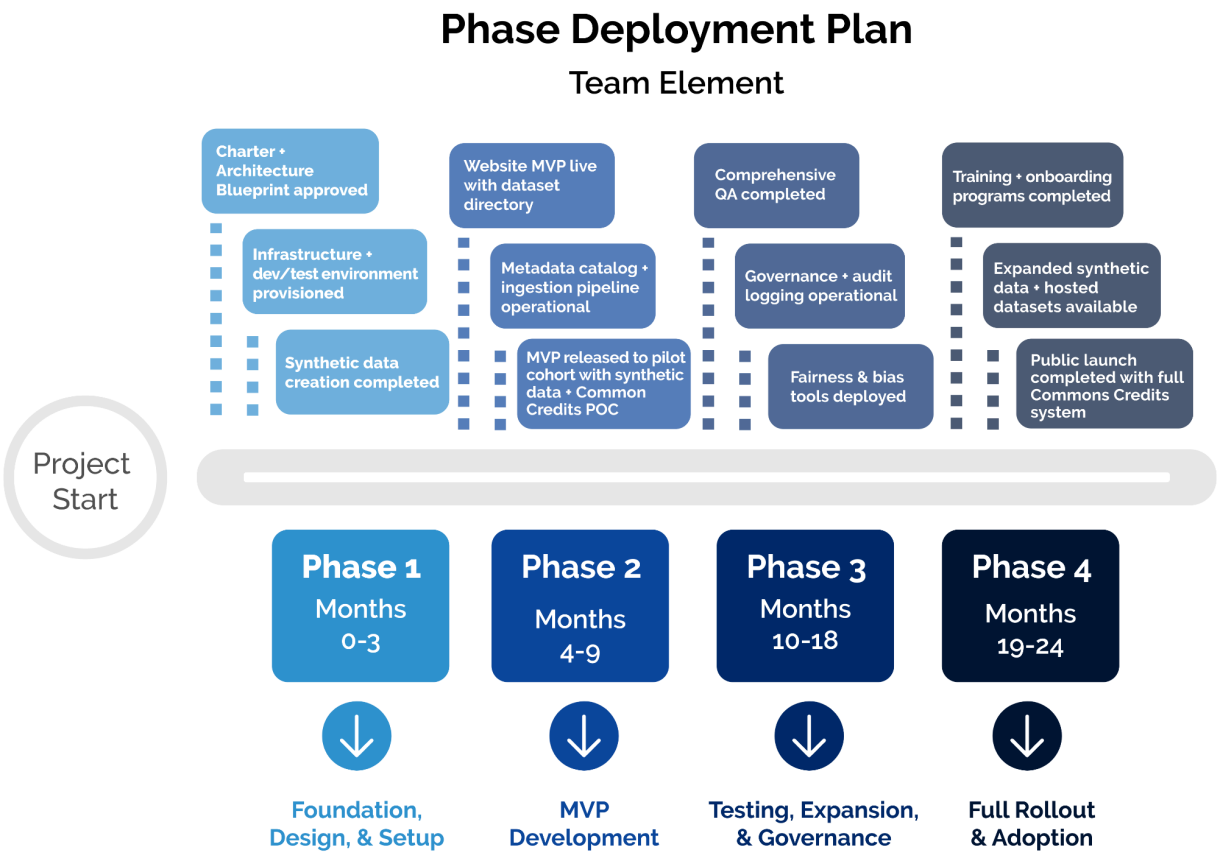
Phase Deployment Plan
Team Element
Charter + Architecture Blueprint approved
Infrastructure + dev/test environment provisioned
Synthetic data creation completed
Website MVP live with dataset directory
Metadata catalog + ingestion pipeline operational
MVP released to pilot cohort with synthetic data + Common Credits POC
Comprehensive QA completed
Governance + audit logging operational
Fairness & bias tools deployed
Training + onboarding programs completed
Expanded synthetic data + hosted datasets available
Public launch completed with full Commons Credits system
Project Start
Phase 1 Months 0-3
Phase 2 Months 4-9
Phase 3 Months 10-18
Phase 4 Months 19-24
Foundation, Design, & Setup
MVP Development
Testing, Expansion, & Governance
Full Rollout & Adoption

2. Technical Approach

2.1 Program Management & Stakeholder Coordination

Overview. Effective program management is the foundation for the success of the DCC. Element, as Prime, will serve as the integrator responsible for aligning scope, schedule, cost, and quality across all workstreams, while JDSAT provides technical program support to ensure synthetic data and AI fairness milestones are fully coordinated with overall program objectives. Together, our integrated team will deliver a disciplined, transparent, and collaborative management approach that ensures MassTech's goals are achieved on time and within budget.

Approach. Element will establish and maintain a comprehensive Integrated Master Project Plan and Timeline, leveraging Scaled Agile Framework (SAFe) controls. Using two-week sprint cycles, retrospectives, and demos, our team will track progress against milestones and rapidly resolve risks, issues, and changes. Element will implement risk registers, Earned Value tracking, and executive-level

dashboards to provide MassTech with transparent, real-time visibility into program health. Monthly executive reviews and weekly program updates will ensure decision-makers remain engaged and informed.

JDSAT will directly support program management by aligning technical deliverables, particularly in synthetic data, AI/ML modeling, and fairness frameworks within the program plan. Their SMEs will participate in working groups, provide updates during stakeholder reviews, and ensure that highly technical milestones integrate seamlessly with governance, compliance, and adoption timelines. This joint governance model will incorporate bi-weekly sprint sessions with stakeholders demos to validate direction, supported by pilot users to capture feedback loops that inform continuous improvement.

Benefits. Through Team Element's agile release train, sprint cycles, and executive dashboards provide decision-makers with real-time visibility into program health, reducing risk and enabling proactive adjustments. Researchers, policymakers, and innovators benefit from stakeholder demos, feedback loops, and governance checkpoints that ensure technical deliverables, like synthetic data pipelines and AI fairness framework, are validated against real-world needs before scaling.

2.2 Platform Architecture & Development

Overview. The Massachusetts AI Hub requires more than a catalog. It needs a platform that creates trust, removes barriers to use, and scales with the Commonwealth's ambition for responsible AI. Our design delivers these outcomes by combining CKAN's strength as a metadata-driven catalog with Dataverse's rigor in research data management, producing a unified discovery and access layer for government and academic users alike.

CKAN. CKAN serves as the central portal where agencies, researchers, and innovators search, filter, and discover datasets. The catalog applies metadata standards such as Dublin Core, Schema.org, and DCAT, ensuring every dataset is consistently described and therefore interoperable across institutions. Faceted search and full-text queries help users quickly locate relevant datasets, while semantic search powered by embeddings allows the discovery of connections that go beyond keywords. This reduces wasted effort and accelerates the pace of research. Metadata will also be structured to support multilingual content. CKAN plugins and Dataverse extensions will allow dataset titles, abstracts, and keywords to be entered and displayed in multiple languages. Where possible, automated translation services can generate initial non-English metadata fields, with curators responsible for steward review to maintain accuracy and context. The search index will be configured to recognize multilingual entries, allowing users to discover datasets using search terms in their own language. Provenance and governance metadata will also respect multilingual requirements. Licensing terms, provenance statements, and usage restrictions will be tagged with language codes to ensure accurate interpretation. These metadata records will remain standard, using DCAT and Dublin Core with multilingual fields encoded according to ISO 639 language codes. By delivering machine-readable formats such as CSV, JSON, and Parquet, the platform supports direct integration with modern analytics and AI workflows.

Dataverse. Dataverse complements CKAN by providing research-grade dataset management. Versioning, DOI assignment, and scholarly citation ensure datasets are reproducible and citable in academic work. Metadata is harvested from Dataverse into CKAN, giving the community a single discovery surface without losing the rigor of formal research curation. Metadata fields in Dataverse will support multilingual content so that dataset descriptions, abstracts, and contributor information can be maintained in multiple languages. Automated translation services may be used to generate initial non-English fields, with stewards responsible for review and approval to preserve accuracy. Provenance records, licensing terms, and related publications will carry language tags using ISO 639 codes, ensuring consistent interpretation across systems. This dual capability means the platform delivers the breadth of a government catalog and the depth required by academic users, while remaining inclusive and accessible to a multilingual research community.

Integration. Integration between CKAN and Dataverse is handled across three layers. First, identity is federated into the Commonwealth's single sign-on system using SAML or OIDC. Roles such as public viewer, researcher, curator, and administrator are consistently mapped across both systems, with multi-factor authentication enforced at the identity provider and step-up MFA for Controlled and Restricted datasets. Element brings direct experience with this challenge, having successfully built and integrated SAML and OIDC identity systems for federal programs at the Centers for Medicare and Medicaid Services, where multiple internal and external users required secure, role-based access. Second, metadata is synchronized through harvesting and mapping, preserving provenance elements such as version notes and related publications. Third, CKAN manages access workflows, brokering requests and steward approvals, and issuing time-limited tokens or workspace policies. When datasets are stored in Dataverse, CKAN requests scoped tokens so that users experience seamless access. The platform as a whole will also ensure broad interoperability with key state and federal systems. Team Element will incorporate standards like HL7 FHIR for healthcare data exchange and ensure seamless integration with the Commonwealth's other cloud or on-prem systems.

Governance and Compliance. Security and governance are embedded by design. Every access event, dataset download, fairness evaluation, and synthetic data generation is logged and retained for at least three years. Immutable logs and dashboards give the Commonwealth clear visibility into how data is used, creating accountability for providers and confidence for regulators. Zero Trust patterns, encryption at rest and in transit, and tiered dataset policies protect sensitive content while ensuring openness where possible.

Application Programming Interface (APIs). APIs and integration points are built for developers and researchers. RESTful APIs and SDKs enable programmatic search, access, and provenance retrieval. Jupyter notebooks and MLFlow pipelines connect directly to the platform. JDSAT contributes technical expertise in this area, drawing on its work with the Department of the Navy's Comprehensive Integrated Analytics program, where the team designed cloud-based analytics platforms, implemented MLOps pipelines, and applied fairness testing frameworks. That past performance demonstrates the ability to integrate programmatic APIs and fairness tools into secure, large-scale environments. In the Massachusetts AI Hub, JDSAT will guide the integration of fairness and bias auditing libraries such as AIF360 and Fairlearn, alongside synthetic data frameworks including SDV and Synthea.

Performance and accessibility are priorities from day one. The portal is WCAG 2.1 compliant and internationalized for multilingual readiness at launch. Search results and page loads target under two seconds through optimized indices, CDN caching, and server-side rendering. The architecture scales horizontally to support thousands of users and datasets while maintaining a 99.9 percent uptime target. These commitments matter because researchers, startups, and agencies need reliable and equitable access to the platform to do their work effectively.

Deployment. Element will lead the deployment of CKAN and Dataverse, integration of SSO and MFA, RBAC design, and implementation of audit dashboards. JDSAT will focus on API design, fairness integration, and data science tooling, ensuring the platform supports responsible AI development. Together, the team brings a proven record of delivering secure, metadata-driven platforms for federal and defense clients, including CMS and Navy Medicine, while applying lessons learned from synthetic data frameworks for the Defense Health Agency. This experience shows the proposed solution is both technically feasible and operationally grounded.

Benefits. With Team Element's integrated CKAN–Dataverse architecture, the Massachusetts AI Hub will gain a platform that is not only technically rigorous but also inclusive, scalable, and trusted. By enforcing FAIR metadata standards, multilingual readiness, and WCAG 2.1 accessibility, the Hub ensures that researchers, agencies, and startups can discover and use datasets seamlessly across disciplines and languages. Programmatic APIs and developer toolkits accelerate innovation by enabling direct integration into Jupyter notebooks, ML pipelines, and analytics workflows. Governance features, like MFA, RBAC, immutable logs, and zero trust, deliver accountability and compliance, giving data providers confidence that their contributions are secure and auditable. With 99.9% uptime and sub-two-second performance targets, the Hub provides the reliability needed for mission-critical research. This foundation empowers Massachusetts to foster responsible AI development at scale, positioning the Commonwealth as a global leader in ethical and equitable innovation.

2.3 Dataset Curation & FAIR Compliance

Overview. For the Massachusetts AI Hub to succeed, its datasets must be trusted, reusable, and auditable. Our approach ensures that every contribution to the Data Commons Collaborative is curated to meet FAIR principles: Findable, Accessible, Interoperable, and Reusable. By combining automated ingestion with human oversight, we create a disciplined intake process that guarantees research-grade datasets aligned with the needs of the Commonwealth's innovation ecosystem.

Ingestion pipelines are orchestrated through Apache Airflow, deployed in Kubernetes at MGHPCC. Airflow provides scheduling, retries, and observability, making it well-suited for managing complex dataset flows. Each dataset submission includes a manifest that defines the handoff contract between provider and platform. The manifest specifies required metadata such as title, description, license, creator, steward, and proposed access tier, along with technical descriptors like file format, schema, and checksums. JSON Schema validation enforces these requirements, while SHACL profiles are available for RDF-based submissions. This contract creates clarity for providers and predictability for stewards.

Once staged, datasets move through a layered process. Structural checks confirm integrity and format compliance. Metadata is mapped to DCAT and Dublin Core and verified against CKAN schema rules to ensure a clean publication. Automated profiling generates statistics on completeness, distributions, and outliers, providing users with transparency on dataset quality. Provenance is captured using PROV standards, documenting source system, processing steps, and steward approvals. Each dataset receives a semantic version and a machine-readable changelog, so updates are clear to both humans and machines.

Governance and Compliance. Safeguarding privacy is central to the process. Pipelines include automated scans for personally identifiable information and sensitive attributes, combining pattern-based rules with trained models. If issues are detected, the workflow pauses for steward review or remediation before proceeding. Service levels ensure that automation provides fast turnaround, with validations complete in minutes, while stewards confirm access tiering and licensing decisions within three business days. Approved datasets are published to CKAN and, when applicable, to Dataverse for research-grade curation. Each step is logged and archived for at least three years, creating an immutable audit trail.

Governance is built into the workflow. Stewards confirm dataset tiering into Open, Controlled, or Restricted categories and approve publication. Restricted datasets may be bound to consent tokens that define the duration and scope of access. Revocations trigger Airflow hooks that remove access and schedule the cleanup of derivatives. This combination of automation and human oversight ensures that compliance is maintained while providers retain confidence in how their data is handled.

Operational Use Cases. Past performance demonstrates our ability to deliver this approach at scale. Element led ingestion and validation pipelines for the Centers for Medicare and Medicaid Services, where we operationalized large-scale measure collection systems that validated, curated, and published datasets under strict compliance rules. That work required integrating schema checks, provenance tagging, and versioning into automated workflows while providing human-in-the-loop review for sensitive measures, experience directly applicable to this program. JDSAT complements this by bringing expertise with the Defense Health Agency's MedSentinel project, where their team designed pipelines that generated and validated synthetic datasets with strong provenance and metadata controls. Together, these examples prove our ability to curate datasets under demanding standards, combining technical automation with compliance-driven governance.

Publication is seamless across the two frameworks. CKAN provides broad cataloging and discovery, while Dataverse ensures that datasets intended for academic research receive DOIs, formal versioning, and citation support. Metadata synchronization between the two ensures a coherent discovery experience, eliminating duplication and inconsistency. Previews in CSV, JSON, and Parquet make data immediately usable, and when conversions are required, the changelog documents any transformations for transparency.

Benefit. Our approach ensures that all datasets are classified, tiered, and secured in line with HIPAA, NIST, and FISMA requirements, while also aligning with MassTech's oversight checkpoints. Researchers

and policymakers will benefit from clear policies for access, use, and stewardship, reducing compliance risk and ensuring ethical AI development. The use of RBAC, MFA, and audit logging strengthens accountability, while explicit ethical terms of use safeguard against misuse. By embedding governance into every stage of data lifecycle management, the Hub will provide a secure, trustworthy foundation that balances innovation with regulatory compliance and public trust. The graphic below shows the provider handoff sequences depicted in the approach above, a portrait size view can found in the appendix.

Provider Handoff Sequence

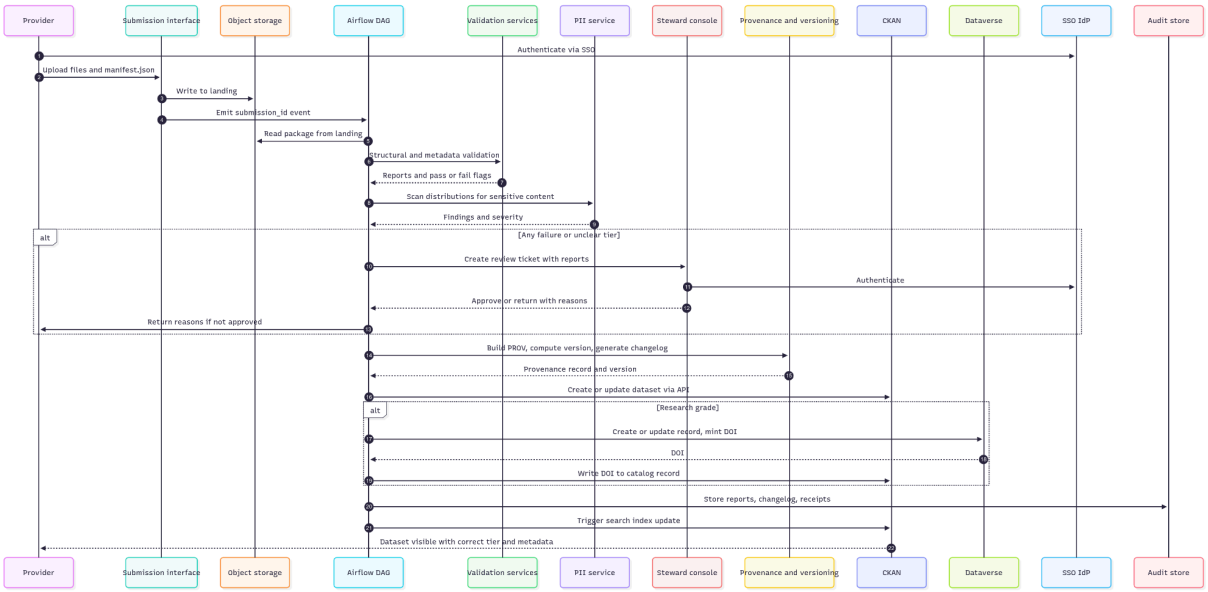

2.4 Synthetic Data Capability

Overview. Our synthetic data capabilities allow agencies to train, test, and validate models in ways that preserve privacy while scaling access to usable data. The result is stronger protections for sensitive information, improved availability of training data in restricted domains, and the ability to simulate rare or edge-case scenarios that will enhance model robustness and operational readiness. Another significant benefit is the ability to begin testing and refining models with synthetic datasets while waiting for production or real-world data to become available, accelerating development timelines and reducing project risk.

Synthetic Data Pipelines. Our pipelines support ingestion of structured, unstructured, time-series, image, and audio data through secure and scalable mechanisms. Automated profiling and quality assessment evaluate data readiness, while de-identification and privacy risk scoring are applied before synthesis. Probabilistic modeling, deep learning methods, and domain-specific simulators support generation. We have general experience with SDV (Synthetic Data Vault), a Python library for tabular

synthetic data, and we operationalize Synthea, which produces realistic synthetic health records in HL7 FHIR format. These tools offer varying strengths: some excel at structured data, others at unstructured or time-series inputs, and some incorporate compliance and privacy validation. Our team understands when to apply each approach and selects the right toolset depending on mission context and data type.

Importantly, our ingestion pipelines enforce FAIR principles and maintain dataset version histories for transparency, reproducibility, and rollback. Metadata governance includes schema alignment with DCAT and Dublin Core standards to promote interoperability across state and federal systems, as well as multilingual metadata support to broaden accessibility.

Our pipelines allow configurable synthesis parameters, enabling agencies to tune fidelity versus privacy tradeoffs, preserve target variables critical to downstream analysis, and adapt datasets to specific mission requirements. Once generated, synthetic datasets are stored with metadata and provenance records that clearly distinguish them from original data. Provenance captures synthesis methods, model parameters, lineage, and version history for auditability. Compute infrastructure provisions scalable CPU and GPU resources optimized for generative modeling and simulation, with containerized pipelines ensuring reproducibility across environments.

Operational Use Cases. The true value of synthetic data lies in how it is applied to mission workflows. In our ambient dictation capability, we integrate Synthea to produce synthetic patient records and dialogue scripts. These are converted into lifelike audio conversations with Amazon Polly, creating a reusable test corpus. The corpus is used to benchmark transcription accuracy, medical entity extraction, and clinical documentation quality, allowing us to validate healthcare AI systems under realistic but privacy-safe conditions.

Our Common Data Processor (CDP) generates synthetic datasets designed to interface with downstream applications for operational readiness, allowing us to test and refine those applications without utilizing sensitive PII or PHI data. Validation ensures schema fidelity, referential consistency, and full coverage of required data fields. The CDP enables us to simulate scenarios, test custom policies, and validate strategies in a secure, privacy-preserving environment.

Governance and Compliance. Governance and compliance are embedded throughout the capability. Stewardship workflows ensure dataset ingestion requests are reviewed and approved before being published into the commons. Sandbox workflows provide controlled environments for dataset ingestion, validation, and governance testing, aligned with MassTech's phased rollout governance checkpoints.

Role-based access controls define permissions for data creators, validators, and consumers, while audit logs capture synthesis activities and access events. Synthetic datasets can be shared more broadly under relaxed governance compared to real data, but accountability is preserved through metadata and provenance tracking. Privacy protections are validated using differential privacy scoring, k-anonymity analysis, and membership inference testing.

Accessibility is prioritized across all interfaces: ingestion portals, metadata search, and documentation comply with WCAG 2.1 standards, ensuring inclusive access for all users. All processes are aligned with HIPAA, state data privacy laws, and other applicable regulations.

Tradeoffs are carefully considered. Higher-fidelity synthesis can deliver greater realism but requires more compute, while lightweight approaches are more efficient but less detailed. Agencies may also choose between a single integrated platform for governance simplicity or specialized tools for domain-specific fidelity. Our approach balances these considerations based on mission priorities, resource constraints, and compliance requirements.

Benefit. By leveraging Team Element's synthetic data pipelines, the Massachusetts AI Hub will gain a powerful capability to accelerate innovation while safeguarding privacy. Agencies, researchers, and startups can train and validate models without exposing sensitive PII or PHI, reducing compliance risk while expanding access to high-value datasets. FAIR-compliant metadata, provenance tracking, and versioning ensure synthetic datasets remain auditable, reproducible, and trustworthy, building confidence across providers and users. Configurable synthesis parameters allow mission owners to balance fidelity and privacy, simulate rare events, and generate domain-specific datasets that improve model robustness. Sandbox workflows provide controlled environments for governance testing, while WCAG 2.1 accessibility and multilingual metadata broaden participation. With differential privacy scoring, k-anonymity, and HIPAA alignment embedded by design, the Hub can scale responsible AI experimentation faster and with less risk.

2.5 AI Fairness & Bias Frameworks

Overview. The AI Fairness and Bias Capability in our solution ensures that all data and AI systems developed within the DCC are equitable, transparent, and aligned with ethical standards. Bias can enter through data, models, or deployment contexts, and if left unaddressed, it can lead to inequitable treatment, reinforce existing disparities, and erode trust in public sector applications. Our approach embeds fairness considerations across the entire AI lifecycle, from data collection and preprocessing of multimodal datasets (text, video, images, audio) to model development, evaluation, and long-term monitoring. By combining governance frameworks with advanced technical methods, we ensure that the DCC delivers AI that serves the public good and contributes to societal benefit.

Technical Methods for Fairness and Bias. We apply both supervised and unsupervised methods to identify, quantify, and mitigate inequities proactively. Supervised fairness testing to evaluate system performance across demographic groups using metrics such as demographic parity and equalized odds. When disparities are found, we apply mitigation techniques such as reweighting or fairness-constrained optimization to improve outcomes. Unsupervised clustering and anomaly detection are used to surface hidden subgroups or patterns of underrepresentation within complex datasets. This is particularly critical in domains like computer vision or natural language processing, where subtle imbalances in training data can have outsized effects on system accuracy and fairness.

To operationalize these practices, we integrate open-source toolkits such as IBM's AIF360 and Microsoft's Fairlearn, which provide standardized methods for bias detection, auditing, and mitigation. These frameworks enable us to establish repeatable, transparent workflows for documenting bias impact and deploying dashboards that monitor fairness metrics throughout the lifecycle. The use of such tools ensures our approach remains state-of-the-art, interoperable, and accountable.

Human-Centric, Trustworthy, and Inclusive Design. Our solution prioritizes human autonomy and dignity by embedding human-in-the-loop safeguards and participatory design practices. Systems are designed to augment rather than replace human decision-making, and user-facing explanations ensure that outputs can be questioned and understood. Trust is reinforced through transparent, auditable pipelines, model documentation, and continuous monitoring for bias and drift. We also ensure that all data and models are secured against unauthorized access and are developed in compliance with applicable state, federal, and international legal standards, reinforcing accountability and public trust.

Inclusivity is addressed by ensuring diverse representation in training data, conducting intersectional fairness analysis, and maintaining feedback loops with affected communities. By aligning technical rigor with social responsibility, the DCC fosters confidence in AI systems and ensures that outcomes reflect a broad range of perspectives and experiences.

Sustainable Innovation. We recognize that fairness extends across both populations today and future generations. Our solution balances innovation with sustainability by embedding energy-efficient practices such as model pruning, quantization, and federated learning to reduce computational costs. Longitudinal fairness monitoring ensures models remain equitable as data distributions and social conditions evolve.

We also acknowledge tradeoffs: fairness constraints may affect predictive accuracy, requiring careful calibration of objectives based on mission context. Similarly, decisions must be made between maintaining a single integrated platform for consistency or tailoring user interfaces to domain-specific needs. At the same time, we position the DCC as a hub for ethical AI innovation by fostering open collaboration with academia, industry, and government; enabling access to high-quality data resources; and supporting rapid experimentation.

Benefit. By embedding fairness checks across the entire AI lifecycle from data ingestion to model monitoring, the Hub ensures that public sector applications serve all communities without reinforcing disparities. Researchers and policymakers will benefit from access to standardized fairness metrics, dashboards, and audit-ready workflows that make bias visible and actionable. Human-in-the-loop safeguards and participatory design practices ensure that outputs are understandable, contestable, and respectful of user autonomy, building public confidence in AI adoption. With inclusive data practices, intersectional analysis, and continuous bias monitoring, the Hub sets a new benchmark for ethical AI innovation.

2.6 Security, Governance & Compliance

Overview. Element will define and enforce data governance policies, implementing encryption, logging, dataset tiering, and ethical terms of use. Our Zero Trust and compliance expertise (StateRAMP, HIPAA, NIST) ensures that security is embedded by design.

Governance and Compliance. Element brings proven expertise in security, compliance, and ethical governance to ensure the Massachusetts AI Hub DCC is trusted, resilient, and regulation-ready from day one. Our approach embeds security by design, aligning with federal and state mandates while enabling responsible AI innovation. We enforce role-based access control and multi-factor authentication to ensure that only authorized users can access the appropriate data, thereby protecting Massachusetts agencies, researchers, and citizens from any unauthorized use. All datasets will be encrypted at rest and in transit to prevent exposure during storage or transfer, and comprehensive audit trails will be retained for three or more years to provide transparency into how data is used, building confidence with regulators and stakeholders. Our archival strategy will use cost effective tiered cloud object storage and can be easily scaled to provide archiving beyond three years if needed.

Element also establishes clear terms of use, ethical guidelines, and dataset tiering, supported by a compliance review board that ensures sensitive or high-impact data is always handled with oversight and accountability. Our solution is designed to meet and exceed Zero Trust principles, StateRAMP, HIPAA, NIST SP 800-53, and Massachusetts privacy laws, assuring that the DCC will remain compliant with evolving mandates. Regular security reviews, penetration testing, and compliance audits further validate protections, keeping the DCC ahead of emerging threats.

Benefits. By embedding Zero Trust principles, StateRAMP, HIPAA, NIST, and Massachusetts privacy requirements into the design, the Hub delivers a resilient platform that protects sensitive data while enabling responsible AI use. Researchers, policymakers, and agencies benefit from role-based access, MFA, and encryption safeguards that prevent unauthorized use and build confidence in data integrity. Immutable audit logs and compliance review boards provide transparency and accountability, giving regulators and stakeholders assurance that datasets are governed ethically and consistently. With proactive penetration testing, ongoing audits, and clearly defined terms of use, the Hub reduces operational risk while fostering innovation.

2.7 Commons Credits Ledger

Overview. We propose a comprehensive Commons Credits System that transforms diverse contributions, monetary, intellectual, and computational, into a unified currency for the Massachusetts AI Hub ecosystem. Our solution creates a transparent, auditable, and scalable platform that seamlessly integrates with major AI platforms while pioneering sustainability incentives. An example implementation could leverage AWS GovCloud's enterprise-grade infrastructure, utilizing services like immudb for immutable transaction recording, DynamoDB for real-time balance management, and Lambda for serverless processing. However, similar architectures could be deployed on Azure, GCP, or hybrid cloud environments based on MassTech's preferences.

Governance Framework. The Commons Credits System will operate under an explicit governance model aligned with MassTech's practices. Oversight committees, including representation from academia, industry, and community stakeholders, will review policies and credit valuation rules quarterly. Transparent reporting cadences (e.g., quarterly dashboards and annual audits) will ensure accountability, while feedback loops and public reporting enhance trust.

Credit Generation and Valuation. Our multi-pathway framework values all contribution types through transparent, market-aligned mechanisms. Monetary contributions follow tiered structures with volume incentives, while grant allocations integrate directly through workflow orchestration services. For instance, an AWS-based implementation would use Step Functions to manage milestone-based releases with special multipliers for priority areas like underserved communities or climate research. Dataset contributions earn credits based on size, quality, uniqueness, and domain relevance, with healthcare or Massachusetts-specific data receiving premium valuations. To strengthen governance, a data classification and tiering model will define categories such as public, controlled, and restricted datasets, each with tailored security and access rules. Human intelligence tasks, including labeling and annotation services, earn credits validated through serverless quality assurance functions, enabling participation through expertise rather than just financial resources.

Universal Exchange Architecture. Our architecture ensures cryptographically verifiable transactions while maintaining high-performance operations. In an AWS deployment, this would utilize the immutable ledger for transaction integrity, with DynamoDB providing millisecond-latency balance queries through global secondary indexes. The system seamlessly converts Commons Credits to platform-specific resources across OpenAI, Anthropic, Hugging Face, and cloud-native ML services. Users select their preferred AI platform or model via a single marketplace-style interface, and the system transparently converts credits into the appropriate tokens or compute units. Guidance tools will also highlight sustainable or cost-efficient choices.

The platform adapter framework could be implemented using API Gateway for RESTful endpoints, with secure credential management through services like AWS Secrets Manager, handling API key rotation. Complex multi-step conversions, such as purchasing external platform credits and updating multiple ledgers atomically, would be orchestrated through workflow services like Step Functions. This design pattern ensures reliability and scalability regardless of the underlying cloud provider selected.

Sustainability Incentives. Our pioneering approach integrates real-time carbon intensity data to incentivize environmentally conscious computing. An example implementation would deploy serverless functions polling APIs like WattTime and Electricity Maps, storing temporal carbon data in time-series databases for efficient analysis. For instance, using AWS infrastructure, Lambda functions would collect this data into Amazon Timestream, though similar architectures using Azure Functions with Azure Time Series Insights or GCP Cloud Functions with BigQuery would achieve comparable results.

Carbon-aware scheduling offers credit discounts for computations during renewable energy peaks, automatically applied through event-driven architectures. Efficiency rewards track when users choose optimized models, creating recurring credits for contributors through custom metrics. Regional

optimization guides users toward low-carbon data centers through globally replicated pricing tables, balancing flexibility with environmental responsibility. Additional sustainability multipliers can reward organizations that adopt model distillation, pruning, or quantization techniques, which reduce computational load and energy use while maintaining accuracy. These practices help ensure that efficiency gains translate directly into both lower costs and reduced carbon impact. These features can be visualized through business intelligence tools like QuickSight, Power BI, or Looker, depending on the chosen platform.

Implementation, Compliance, and Security. Our phased deployment approach ensures the Commons Credits System is delivered securely, transparently, and with the flexibility to expand as the Massachusetts AI Hub grows. Phase 1 establishes the core ledger for immutable transaction recording. Every issuance, transfer, and redemption will be written to a tamper-proof store, creating a permanent and auditable history. Phase 2 integrates the ledger with major AI platforms through API management services, allowing credits to be applied directly to compute, datasets, and analytic workflows. Phase 3 adds sustainability features, including carbon-aware scheduling and time-series analytics, enabling Massachusetts to lead in both equitable and environmentally conscious AI adoption.

The system aligns directly to StateRAMP control families and maps security functions to specific requirements. Access Control is implemented in accordance with AC-2 Account Management, AC-3 Access Enforcement, and AC-6 Least Privilege. Role-based policies ensure that public users, researchers, stewards, and administrators only receive access consistent with their role. Identification and Authentication integrates with the Commonwealth's single sign-on service using SAML or OIDC, meeting IA-2 Identification and Authentication, IA-5 Authenticator Management, and IA-8 Identification and Authentication for Non-Organizational Users. Multi-factor authentication is enforced for Controlled and Restricted datasets, addressing IA-2(1) and IA-2(12). If required, we can also deploy a dedicated identity provider to support sandbox or external tenants.

In addition to role-based access control, the Commons Credits System incorporates tokenized, revocable, and time-bound consent mechanisms. Every dataset or resource request is governed by a cryptographic token that defines scope, duration, and purpose. Tokens can be revoked instantly by data stewards or contributors, ensuring that consent remains dynamic and adaptable to governance policies. Time-bound consent prevents open-ended access to sensitive datasets, while consent events are recorded in the ledger for traceability. This ensures compliance with privacy and ethical use requirements while aligning with HIPAA, FISMA, and NIST frameworks.

Audit and Accountability is supported by immutable event logs meeting AU-2 Audit Events, AU-6 Audit Review, and AU-11 Audit Record Retention. Every ledger transaction, API call, and administrative action is logged and retained for at least three years. Audit logs are cryptographically protected in line with AU-9 Protection of Audit Information and integrated into the state's SIEM for correlation and alerting. System and Communications Protection is enforced through encryption in transit and at rest with SC-12 Cryptographic Key Establishment and Management, SC-13 Cryptographic Protection, and SC-28

Protection of Information at Rest. APIs are secured through SC-7 Boundary Protection, SC-23 Session Authenticity, and rate limiting at the API gateway.

Configuration Management is achieved through CM-2 Baseline Configuration, CM-3 Configuration Change Control, and CM-6 Configuration Settings. Infrastructure-as-code templates enforce consistent configuration across environments and ensure that every change is documented and approved. Incident Response capabilities align to IR-4 Incident Handling and IR-5 Incident Monitoring, with ledger anomalies and unauthorized access attempts automatically escalated to state monitoring and response teams.

In cloud deployments, these controls map to managed services. On AWS, KMS enforces SC-12 and SC-13 for encryption, WAF addresses SC-7 for API protection, and IAM Identity Center integrates with state SSO to meet IA-2 and IA-8. Azure provides parallel functions through Key Vault, Active Directory, and Sentinel, while Google Cloud achieves the same with Cloud KMS, Cloud Identity, and Cloud Logging. This ensures Massachusetts avoids vendor lock-in while remaining StateRAMP compliant.

Accessibility and governance requirements are also addressed. User-facing services conform to WCAG 2.1 Level AA standards, multilingual support is included from launch, and sandbox environments allow governance policies to be tested under CM-4 Security Impact Analysis without impacting production. All usage is bound by ethical terms of service, meeting SA-9 External System Services and PL-4 Rules of Behavior.

By mapping controls directly to the StateRAMP framework, the Commons Credits System delivers transparency, accountability, and public trust. The immutable ledger provides a permanent audit trail, while the modular cloud-native design scales with adoption. The system allows Massachusetts to unify diverse contributions and resource consumption through a secure credit model that is auditable, accessible, and compliant with state security standards.

Benefits. Our Commons Credits System positions Massachusetts as the leader in equitable and sustainable AI innovation by eliminating access barriers between researchers and enterprise-grade AI resources. The cloud-native architecture, whether deployed on AWS, Azure, or GCP, ensures consistent performance scaling from hundreds to hundreds of thousands of users. The immutable ledger design guarantees transparency essential for public trust, while creating the first systematic rewards for carbon-efficient AI development at scale.

By unifying resource management across multiple AI platforms through a single currency, we dramatically simplify operations for both individual researchers and large organizations. The system values diverse contributions from datasets to compute, creating an inclusive ecosystem where universities, communities, and businesses naturally collaborate. Market-responsive exchange rates maintained in high-performance databases ensure long-term financial sustainability, while the cloud-agnostic design protects MassTech's investment by avoiding vendor lock-in. This flexibility, combined with our proven expertise in cloud architectures, makes our team the ideal partner for delivering this transformative capability to the Commonwealth's AI community.

2.8 Testing, Rollout & Adoption

Overview. Team Element proposes a comprehensive, audit-ready testing and adoption framework for the Massachusetts AI Hub, grounded in our proven integration testing practices from the Quality Payment Program (QPP) at CMS. Our approach ensures that the DCC is technically sound, ethically grounded, and fully embraced by its intended users through rigorous quality assurance, robust training, thorough documentation, and a structured public launch strategy.

Testing Framework. Our commercial-grade Test Scenario Management (TSM) application enables end-to-end QA across all platform functions, including eligibility, data submission, performance feedback, targeted reviews, and exception workflows. With over 1,000 synthetic test scenarios designed to replicate real-world conditions while maintaining data privacy, we ensure functional, performance, and security testing of all platform components. This includes synthetic data validation with reproducible outcomes, UI-to-API consistency checks, automated validations (25+ active checks), and continuous monitoring across production and test environments. To ensure consistency across user interfaces and backend systems, we employ a screen scraping process that validates UI behavior against API responses in lower environments. This method confirms that synthetic test data renders accurately across modalities, critical for fairness dashboards, model transparency, and user trust. Ensuring DCC's platform uptime and operational stability, early detection of performance drift or feature regressions, and real-time visibility into system health.

Category	Measurable Outcomes	Responsibility
Testing Framework	≥95% defect detection prior to production release; uptime maintained at ≥99.9%; quarterly penetration testing reports delivered.	Element leads QA governance and monitoring; JDSAT supports synthetic data validation and fairness testing.
Training & Adoption	≥80% of target users trained before launch; ≥90% satisfaction rating in post-training surveys.	Element leads training delivery and adoption strategy; JDSAT provides technical content for synthetic data and fairness modules.
Documentation & Support Materials	Documentation updated within 10 business days of any system release; ≥90% of user-reported documentation needs addressed in feedback cycles.	Element manages the documentation lifecycle; JDSAT contributes technical specifications for AI/ML, fairness, and synthetic data modules.
Public Launch & Adoption	≥90% of pilot feedback incorporated before full launch; ≥75% of intended users onboarded within the first 30 days.	Element leads rollout coordination and stakeholder engagement; JDSAT provides technical support for data/scenario demos and validation.

Training and Adoption. To ensure successful onboarding, sustained engagement, and long-term usability of the Massachusetts AI Hub, Element and JDSAT will deliver a comprehensive training and adoption strategy. Drawing on our experience supporting federal programs, we will provide tailored training for both end users and system administrators. This includes virtual workshops, sandbox environments for safe experimentation and training, role-specific onboarding materials, and access to ongoing support channels such as help desks, office hours, and community forums. All training content will be designed to meet FAIR principles , adhere to Section 508 accessibility standards, WCAG 2.1 standards with multilingual accessibility from launch, and be audit-ready for compliance and governance review.

Documentation and Support Material. We will deliver comprehensive, version-controlled documentation, including detailed user guides, system administration manuals, and a searchable knowledge base with FAQs and fact sheets. Documentation will align with the RFP's deliverables table, including the Training Materials & User Guides, Commons Credits Ledger, and Governance & Security Package. Each resource will be version-controlled, regularly updated, and structured to support self-service learning and operational continuity.

Public Launch and Adoption. Our phased rollout strategy ensures a smooth and successful public launch. Activities include targeted demos, pilot programs, iterative feedback loops, launch-day monitoring, and post-launch support. Stakeholder walkthroughs will be tailored to researchers, data stewards, and policy stakeholders, ensuring that platform capabilities are clearly understood and aligned with user needs. Pilot feedback will be captured through structured surveys and engagement sessions, informing refinements before full deployment.

Benefits. Comprehensive QA built on 1,000+ synthetic test scenarios and continuous monitoring ensures defects are caught before release, uptime exceeds 99.9%, and fairness dashboards perform reliably. Researchers and agencies gain assurance that the Hub is resilient, auditable, and privacy-preserving, while tailored training, multilingual documentation, and sandbox environments equip users to adopt the platform quickly and effectively. Structured demos, pilot feedback loops, and phased rollout minimize risk and align capabilities with stakeholder needs, ensuring smooth onboarding and early user satisfaction. By combining rigorous quality controls with inclusive adoption strategies, the Hub will achieve rapid, sustainable engagement.

3. Deliverables

Deliverables are the tangible outputs that will demonstrate progress, assure accountability, and provide value to the Massachusetts AI Hub and its stakeholders. Our team is committed to producing high-quality deliverables that are actionable, transparent, and aligned with FAIR principles, governance frameworks, and stakeholder needs. Each deliverable will be created in accessible formats with supporting documentation to ensure long-term sustainability, usability, and integration into the DCC.

Team Element will manage deliverable quality and timeliness through structured review cycles, agile sprint ceremonies, and formal sign-off processes. Deliverables will be traceable to program milestones, ensuring MassTech receives clear visibility into progress and impact at every stage. The table below depicts a structured list of deliverables that demonstrates how Team Element will meet and exceed the RFP requirements.

Deliverable	Description	Responsible Party
Project Plan & Timeline	Integrated Master Schedule with milestones, dependencies, and risk registers.	Element (Lead), JDSAT (Support)
Architecture Documentation	Detailed technical architecture for metadata systems, APIs, security, and integration.	Element (Lead), JDSAT (Support)
Metadata Catalog & Data Management System	FAIR-compliant catalog using CKAN/Dataverse frameworks, with ingestion pipelines and provenance tracking.	Element (Lead), JDSAT (Support)
Synthetic Data Tools & Outputs	Synthetic datasets generated using Synthea, SDV, CARLA, etc., with validation pipelines and privacy safeguards.	JDSAT (Lead), Element (Governance Review)
Fairness & Bias Reporting Dashboards	Dashboards implementing AIF360 and Fairlearn to monitor and mitigate bias in datasets and models.	JDSAT (Lead), Element (Support)
Commons Credits Ledger	Ledger system for tracking dataset usage credits, with governance and reporting mechanisms.	Shared Responsibility
Governance & Security Package	Policies, controls, and compliance frameworks including RBAC, MFA, encryption, and audit trails.	Element (Lead)
AI Hub Website with Dataset & Login Access	Public-facing AI Hub portal providing secure dataset access, account management, and user dashboards.	Element (Lead – platform & access), JDSAT (Support – dataset integration)
Training Materials & User Guides	Comprehensive training, onboarding, and adoption toolkits for stakeholders and administrators.	Element (Lead), JDSAT (Technical Inputs)
Pilot User Feedback Report	Findings from pilot users, including lessons learned, feedback, and recommendations for platform improvement.	Element (Lead), JDSAT (Support)

By balancing Element's leadership in governance, program management, and adoption with JDSAT's technical depth in synthetic data and AI fairness, we will ensure that each deliverable drives transparency, trust, and meaningful outcomes for the Massachusetts AI Hub.

4. Past Performance & Differentiators

1. Element Past Performance: Quality Payment Program (QPP) Clinician Interface & Scoring (CIS) II
Client: Department of Health and Human Services, Centers for Medicare & Medicaid Services
Contract Number: GS35F372DA
POC: Mindy Riley, mindy.riley@cms.hhs.gov, 443-340-6352
Period of Performance: May 2024 - May 2029
Contract Value: Elements portion ($4M)

Project Summary: Since 2009, Element supported the Centers for Medicare & Medicaid Services (CMS) Centers for Clinical Standards and Quality (CCSQ). Under the Quality Payment Program (QPP) Clinician Interface and Scoring (CIS) II contract, we served as the integration team, collaborating with 200+ contractors to modernize quality reporting. QPP is a nationwide initiative advancing value-based care by improving outcomes, quality, and efficiency, which are directly relevant to building the Massachusetts AI Hub DCC.

Challenge. CMS's manual testing processes limited agility and accuracy in validating quality measure calculations. Inefficiencies drove higher costs, increased risk of production defects, and delayed feature releases. Incomplete test scenarios and unstable data further undermined scoring reliability, echoing the same issues the DCC seeks to address around trustworthy data, governance, and compliance.

Solution Approach. Element introduced automation and analytics to transform testing and data validation, delivering capabilities directly aligned with the AI Hub RFP.

Automated Testing & QA Frameworks. Implemented automated scoring integration tests, API-based submissions, and data-driven analytics to detect defects early and improve validation accuracy.

Test Scenario Management (TSM). Built a centralized repository for synthetic and real test data, enabling end-to-end testing across multiple program components. TSM ensured provenance, lineage tracking, and reproducibility, paralleling RFP requirements for metadata and governance.

Synthetic Data & Privacy. Leveraged Synthea to generate realistic, privacy-preserving health records in HL7 FHIR, enabling safe testing and scenario coverage without exposing PHI/PII.

AI-Enabled Innovation. Piloted an AI-powered chatbot for QPP stakeholders, testing fairness, explainability, and reliability of LLMs against vetted data sources, a direct precursor to the AI Hub's fairness and responsible AI objectives.

Agile Delivery & Governance. Used SAFe Agile, RBAC, audit logging, and compliance alignment with HIPAA and NIST standards to ensure secure, scalable delivery.

Outcomes & Results:

- Reduced production defects by 85%, improving reliability and user trust.
- Accelerated feature deployment, enabling faster delivery of program updates.
- Expanded test coverage with 1,000+ scenarios, reducing risk of edge-case failures.
- Freed technical staff from repetitive tasks, increasing focus on mission-critical innovation.
- Improved scoring accuracy and transparency, driving provider participation and confidence.
- Enabled cost savings that allowed reinvestment into higher-value features.

Relevance to Massachusetts AI Hub DCC:

This past performance demonstrates Element's ability to deliver a scalable, FAIR-aligned, synthetic-data-enabled testing platform with strong governance and compliance. Our work at CMS proves we can:

- Build metadata-driven pipelines with provenance and versioning.
- Implement synthetic data generation with privacy safeguards.
- Deliver automated QA/testing aligned with governance checkpoints.
- Integrate AI fairness and explainability into production systems.

2. JDSAT Past Performance: Comprehensive Integrated Analytics (CIC)

Client: Department of the Navy, Bureau of Medicine and Surgery
POC: CAPT Tracy Krauss, tracy.r.krauss@health.mil, 703-681-9264
Contract Number: N0018922F3052
Period of Performance: September 2022 - September 2027
Contract Value: $29,538,127
Project Summary: JDSAT delivered professional services to improve data and analytic maturity across Navy Medicine, supporting the health and readiness of Sailors and Marines while improving operational outcomes. Our team established a comprehensive analytics framework that transformed operational decision-making capabilities through advanced data governance, cloud architecture, and AI/ML implementation.

A cornerstone of our work involved developing Navy Medicine's Data Transformation Strategy across four pillars: Culture and Talent, Data Governance, Architecture and Tools, and Business Process Integration. JDSAT established the governance infrastructure by drafting the NAVMED Data Governance Board Charter, Data Management Instruction, and Data Sharing Standard Operating Procedures, which significantly improved a common data repository, enhanced data quality, and increased user accessibility.

Our team architected and implemented a cloud-based analytics environment, developing the solution architecture for both classified and unclassified platforms. We deployed multiple custom applications to the Medical Information Platform (MIP), establishing a secure foundation for advanced analytics. The deployment of containerized analytic tools prepared Navy Medicine for transitioning to cloud-based operations, improving scalability and accessibility.

JDSAT implemented a functional MLOps pipeline in a secure cloud environment using Bureau of Medicine Information System (BUMIS) data. This pipeline enables consistent model development, validation, and deployment, ensuring that updates are reproducible and scalable. The infrastructure leverages both traditional machine learning techniques (e.g., clustering algorithms) and modern neural networks to analyze Navy Medicine message traffic. Anomaly detection methods were applied to extract infrequent but crucial message types, related to personnel readiness and equipment or supply needs, fed into downstream

analytical tools for deeper assessment. This MLOps foundation ensures that AI/ML capabilities are reliable, transparent, and aligned with operational requirements.

To reduce reliance on sensitive operational data, JDSAT's Synthetic Common Data Processor (CDP) supports the creation of realistic synthetic datasets that replicate key attributes of billets, personnel, specialties, substitutions, and units. These synthetic datasets allow planners and developers to test models, validate assignment strategies, and run large-scale simulations without exposing actual personnel records. By enabling experimentation with representative but privacy-preserving data, this capability accelerates model refinement, enhances resilience, and safeguards sensitive information while still providing comprehensive coverage of real-world conditions.

Through these comprehensive analytics solutions, JDSAT transformed Navy Medicine's ability to leverage data as a strategic asset, creating systems designed to augment human decision-making with transparent, auditable pipelines, model documentation, and continuous monitoring for bias and drift capabilities directly applicable to the Massachusetts AI Hub's Data Commons Collaborative initiative.

Outcomes & Results:

- Established Navy's first comprehensive Data Transformation Strategy, establishing an architecture that aligned 12 disparate data systems.
- Increased data accessibility from 0 to 20 operational decisionmakers for critical medical readiness applications and data visualizations.
- Shortened decision-making cycle time from 6-8 hours to 5-7 minutes with optimization models linked to automated, transformed personnel data.
- Deployed containerized applications to secure cloud environments with full RMF compliance.
- Developed synthetic, privacy-preserving datasets spanning billets, personnel, specialties, and units, enabling large-scale testing and simulation without exposing sensitive records.

Relevance to Massachusetts AI Hub DCC:

This past performance demonstrates JDSAT's ability to deliver enterprise-scale data transformation and AI enablement with robust governance frameworks. Our work with Navy Medicine proves we can:

- Design and implement comprehensive data governance structures with clear policies and procedures.
- Deploy secure, containerized applications in classified and sensitive environments.
- Create synthetic datasets that enable risk-free simulation and validation.
- Develop cloud-based ML pipelines
- Transform siloed data assets into unified, accessible resources for diverse stakeholders.
- Balance innovation with security and compliance in government environments.

3. JDSAT Past Performance: MedSentinel Synthetic Data Framework

Client: Defense Health Agency (DHA)
POC: LT Adam Braly, adam.m.braly.mil@us.navy.mil, 703-589-5742
Project: MedSentinel Expeditionary Medical C4I
Period of Performance: Sep 2024- Jan 2027
Contract Value: $1,299,997
Project Summary: In conjunction with the MedSentinel Edge Computing project, JDSAT is developing AI/ML capabilities for deployment on MedSentinel. One of these, MediScribe, provides ambient listening at the tactical edge to generate clinician documentation in real time. To support this, JDSAT built a MediScribe prototype with a synthetic data generation framework that enables rigorous testing and validation of AI-powered clinical documentation systems. The framework leverages Synthea to produce synthetic patient records and dialogue scripts, which are then converted into lifelike audio conversations via Amazon Polly, creating a reusable test corpus. This approach provides high-fidelity synthetic data while eliminating privacy concerns associated with protected health information. The resulting synthetic datasets serve as benchmarks for performance evaluation, enabling rigorous measurement of word error rate (WER) for transcription as well as precision, recall, and F1 scores for clinical entity extraction.

JDSAT is now building a production version of this capability designed for operational deployment on MedSentinel. This version incorporates role-based access controls to govern permissions for data creators, validators, and consumers, while also embedding metadata and provenance records that document synthesis methods, model parameters, and data lineage. In parallel, JDSAT is evaluating lightweight, offline AI models tailored for the tactical edge, where compute, power, and connectivity are often limited. This includes models optimized to run on CPU-only environments (when GPU is not available), along with distillation and quantization techniques that preserve accuracy while reducing computational overhead. Together, these features ensure synthetic data pipelines and AI/ML applications like MediScribe can be securely managed, fully auditable, and operationally effective—even in constrained or disconnected environments, while seamlessly integrating with the broader containerized AI pipeline.

The containerized AI pipeline provides reproducibility across environments and scalable CPU and GPU resources where available. It supports the generation of diverse data modalities—including structured, unstructured, time-series, image, and audio—establishing a strong foundation for secure and ethical AI development in sensitive healthcare contexts. Together with the MediScribe synthetic data framework, this infrastructure demonstrates JDSAT's ability to deliver privacy-preserving, high-fidelity training data and optimized edge-ready models that enable advanced AI capabilities while meeting stringent compliance and operational requirements.

Outcomes & Results:

- Reduced development costs by eliminating reliance on sensitive production data.
- Accelerated AI model development cycles through continuous synthetic data generation pipelines.
- Implemented role-based access controls, reducing unauthorized data access incidents to zero.
- Established comprehensive provenance tracking for all synthetic assets, ensuring full auditability.
- Created high-fidelity synthetic audio datasets to benchmark transcription and clinical entity extraction models.
- Eliminated PHI exposure risks by replacing sensitive patient data with synthetic alternatives.
- Enabled deployment of CPU/GPU-flexible, resource-efficient models for tactical edge environments.

Relevance to Massachusetts AI Hub DCC:

This past performance demonstrates JDSAT's expertise in building production-grade synthetic data generation frameworks that enable trustworthy AI development. Our work with DHA proves we can:

- Implement Synthea and other open-source tools to generate domain-specific synthetic data at scale.
- Create multimodal synthetic datasets (text, audio, structured, unstructured) with high fidelity to real-world data.
- Design comprehensive metadata and provenance systems that track synthetic data lineage.
- Implement fairness testing and bias detection in synthetic data generation processes.
- Ensure privacy and regulatory compliance while enabling innovation through synthetic alternatives to sensitive data.
- Incorporate sustainability considerations by developing models that can run flexibly on CPU or GPU depending on resource requirements, and by applying techniques such as model distillation and quantization to reduce compute and power demands while preserving accuracy.

5. Risk Mitigation

Team Element will apply a proactive risk management approach aligned with industry best practices (PMI, NIST RMF) to ensure the successful delivery of the DCC. Our plan identifies potential risks early, assesses their likelihood and impact, and implements clear mitigation strategies with assigned ownership. Risks will be tracked through a living risk register, reviewed during bi-weekly program meetings, and reported in monthly status updates to MassTech, as depicted in the table below

Risk Register

Risk	Impact	Likelihood	Migration Strategy	Owner
Technical Integration Challenges – Difficulty integrating metadata tools (CKAN/Dataverse) with APIs and ledger	Delays in platform readiness; increased rework	Medium	Use proven open-data integration SMEs; conduct phased prototyping and integration testing	Element (Lead), JDSAT (Support)
FAIR Compliance Gaps – Datasets fail to meet FAIR standards	Reduced usability and credibility of platform	Medium	Implement automated metadata validation; JDSAT performs quality checks; Element enforces governance rules	Shared
Synthetic Data Validation Issues – Synthetic datasets lack fidelity or violate privacy	Risk of non-compliance, reduced trust in platform	High	JDSAT to implement validation pipelines (Synthea, SDV, CARLA); apply differential privacy; Element ensures compliance review	JDSAT (Lead)
Bias and Fairness Gaps – AI fairness frameworks (AIF360, Fairlearn) not fully adopted	Reputational risk; platform undermines equity mission	Medium	Deploy fairness dashboards; run pilot audits; partner with academic reviewers; Element embeds fairness into governance	JDSAT (Lead), Element (Support)
Stakeholder Adoption Resistance – Users slow to adopt platform or resist changes	Reduced utilization of DCC	Medium	Element leads stakeholder engagement, HCD training, and onboarding workshops; JDSAT provides user-friendly analytic dashboards	Element (Lead)
Security & Governance Non-Compliance – Failure to meet RBAC, MFA, encryption requirements	Data breach risk; loss of trust; possible legal impact	Low	Implement Zero Trust model; enforce MFA and dataset tiering; conduct regular security audits	Element (Lead)
Budget/Schedule Overruns – Delays in delivery or underestimated costs	Strain on program resources	Medium	Element applies agile project controls; bi-weekly sprint reviews; shared Earned Value tracking	Element (Lead), JDSAT (Support)

Our approach emphasizes early detection and proactive mitigation. Risks are logged in a centralized risk register managed by Element, with each risk assigned a clear owner (Element, JDSAT, or shared). Risks are scored for likelihood and impact, and mitigation plans are revisited during sprint retrospectives.

Element will lead monthly governance reviews with MassTech to report on risk status, identify new risks, and track mitigation progress. JDSAT will provide technical risk input, particularly for synthetic data and fairness frameworks. This joint governance model ensures risks are managed holistically across program management, technical integration, compliance, and adoption dimensions.

6. Team Structure

Team Element is designed to balance strategic leadership, technical expertise, and execution capacity to deliver the DCC successfully. Our team structure ensures senior-level oversight from seasoned program leaders, deep subject matter expertise from senior architects and data scientists, and execution support from mid- and junior-level staff to ensure scalability and cost-effectiveness.

Role Title	Role Description
Shared Services Team:	
Program Director	Program leadership, PMO, governance integration
Solutions Architect	Hybrid cloud architecture, technology selection (CKAN, etc.), platform integration, API strategy, scalability planning
Governance Lead	Data governance, compliance, policy oversight, FAIR compliance, metadata management, curation
Security Architect	Security framework, RBAC, Zero Trust, compliance
Scrum Master	Requirements gathering, stakeholder support
DevOps Engineer	CI/CD pipelines, cloud deployment, containerization, infrastructure automation
Training & Adoption Lead	User training, onboarding, adoption support
Portal Team:	
UI/UX Designer	Accessibility Compliance
Full-Stack Developer (Sr)	Full-stack development, API integration, system architecture, code optimization
Full-Stack Developer (Jr)	Full-stack development, API integration, system architecture, code optimization
Platform Architecture & Development Team:	
System Engineer (Sr)	Platform engineering, integration, DevSecOps, Q and A integration
Software Engineer (Jr)	Application coding, bug fixing, unit testing, integration testing, and regression testing, documentation support
Data Engineer	data pipelines, data ingestion, format conversion, data modeling, versioning support
Full-Stack Developer (Jr)	Full-stack development, API integration, system architecture, code optimization
Data Compliance & Infrastructure Team:	
Chief Data Scientist	Synthetic data, AI/ML strategy, technical leadership
Data Architect	Data platform design, schema, scalability
Data Analyst	Data profiling, statistical analysis, reporting, data quality monitoring
AI Engineer	ML modeling, fairness frameworks (AIF360, Fairlearn)
Synthetic Data Engineer	Synthea/SDV pipelines, validation, privacy
Software Engineer (Sr)	Application coding, bug fixing, unit testing, integration testing, and regression testing, documentation support

The Shared Services Team will support the other three teams emphasizing integration and collaboration ensuring continuity. Element and JDSAT staff will work as a unified team, participating in joint governance reviews, agile sprint ceremonies, and cross-functional working groups. Team Element has not identified specific key personnel at this time. However, we will work with MassTech to designate and commit qualified staff upon award to ensure alignment with program needs and priorities. This approach

ensures seamless coordination across program management, technical development, dataset curation, synthetic data validation, fairness frameworks, and adoption activities.

7. Conclusion

By selecting Team Element, the Massachusetts AI Hub gains a partner dedicated to delivering a scalable, secure, and ethical Data Commons Collaborative. Our proven expertise in metadata management, secure access, AI fairness, and synthetic data will provide Massachusetts AI Hub stakeholders with a trusted platform for AI innovation, fairness, and collaboration. We are committed to an adaptive, human-centric, trustworthy, inclusive, innovative, and sustainable DCC that aligns AI development with the public interest and respects the rights of the Commonwealth's citizens. Our realistic and phased approach ensures compliance with state and federal regulations, drives early adoption, and guarantees the long-term sustainability of this vital initiative. This commitment will firmly establish Massachusetts as a national leader in responsible AI data infrastructure, unlocking economic opportunities and preparing the state's workforce for the future of AI.

Appendix

Provider Handoff Sequences

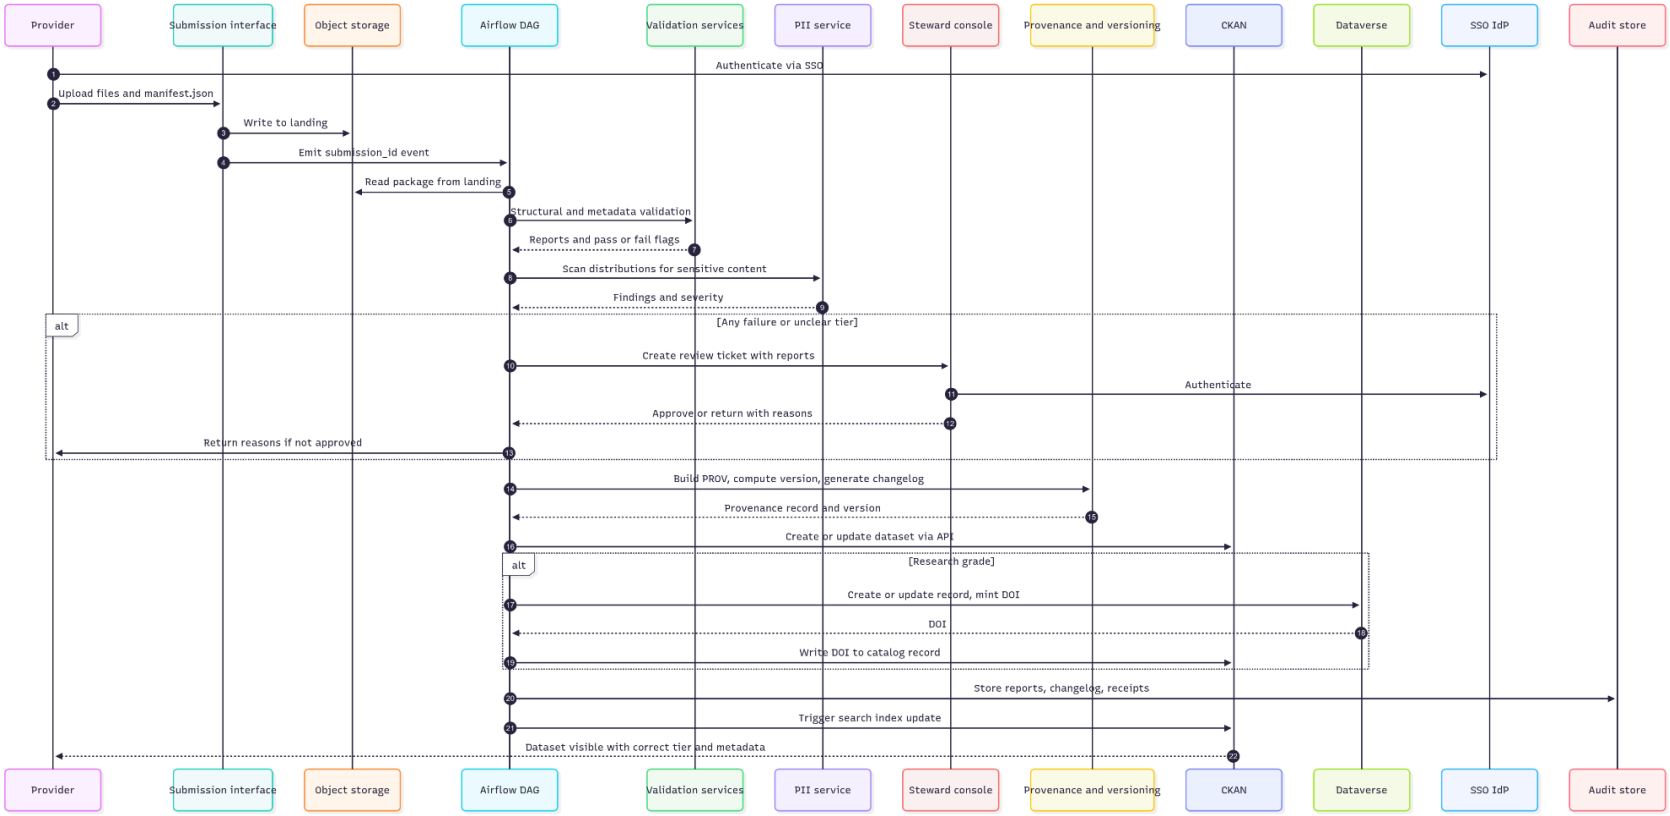

Component and Integration Map

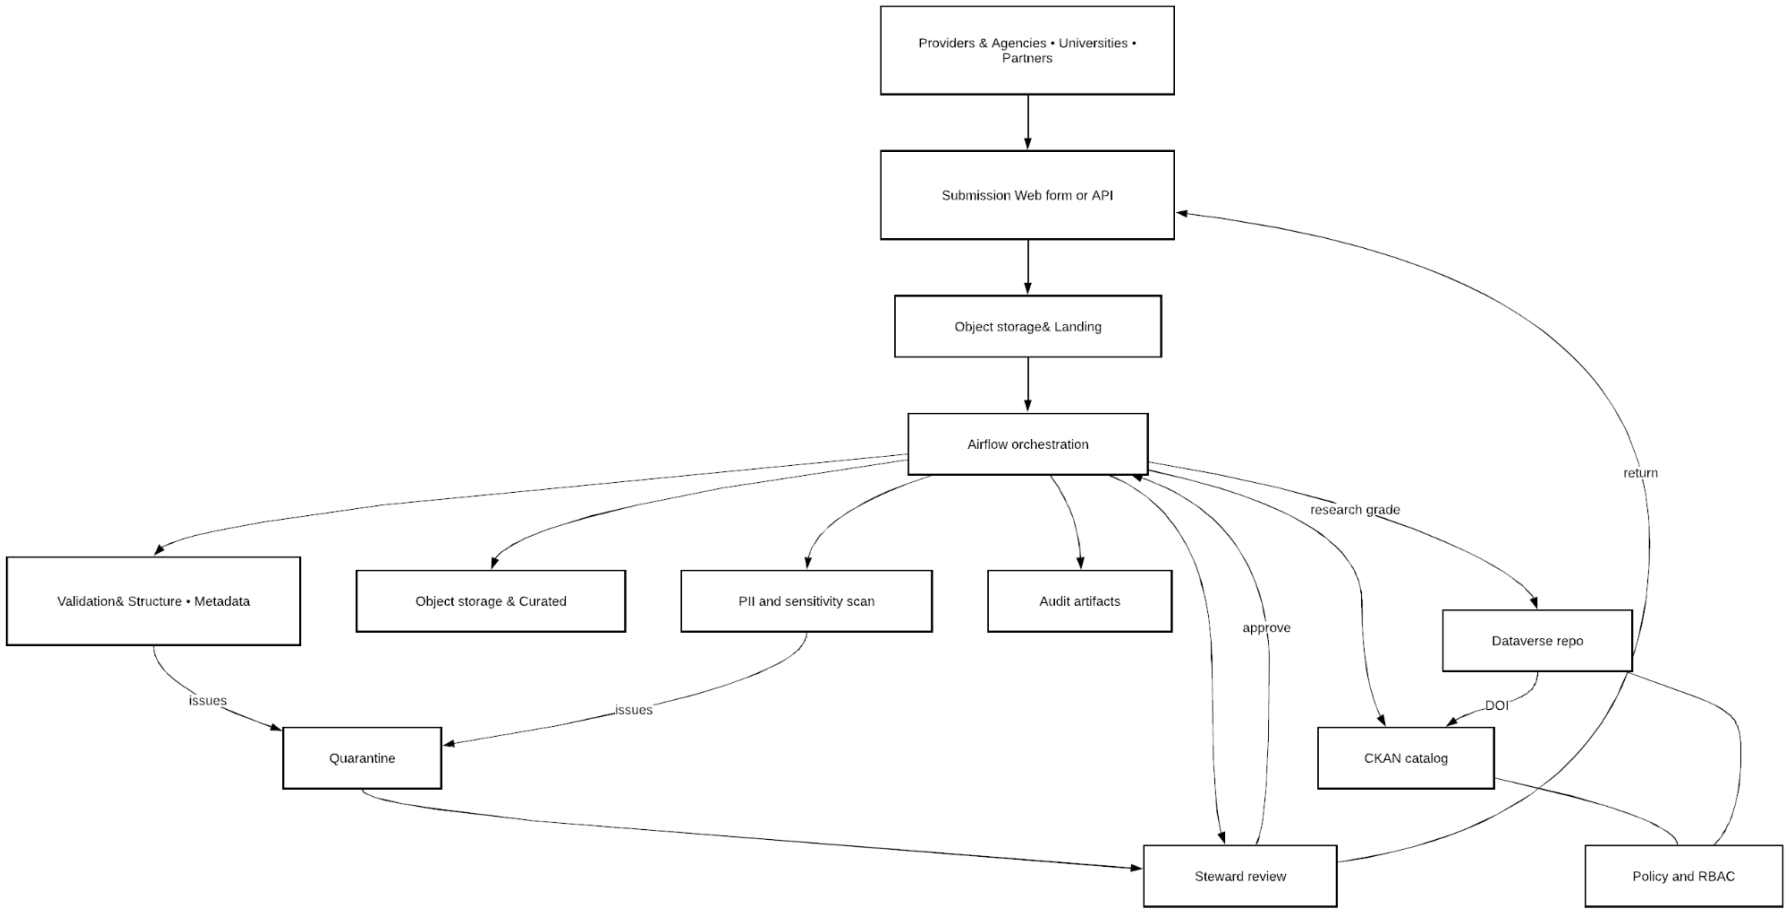

Budgeting

Applicant Information

Applicant: Element Solutions LLC	Procurement No.: 2026-JAII-01
Budget Period: 12/2025 - 12/2027 (2 years)	Title of Proposed Project: AI Hub DCC
Address: 14 Ridge Square NW, Suite 300, Washington, DC 20016	
	MTC Funding Requested:
	Total Project Cost:

LEGEND:

Formula cell

Data entry cell

This section is for post-contract invoicing only. Do not complete with proposal.

Cost Elements			Total Project	Invoice 1 Period: MM/DD/YYYY-MM/DD/YYYY		Invoice 2 Period: MM/DD/YYYY-MM/DD/YYYY		Invoice 3 Period: MM/DD/YYYY-MM/DD/YYYY		Invoice 4 Period: MM/DD/YYYY-MM/DD/YYYY		Cumulative Period: MM/DD/YYYY-MM/DD/YYYY	
					Amount		Amount		Amount		Amount	Expended	Remaining
I. Direct Labor													
Element Solutions LLC	*HRS*	*Rate*		*HRS*		*HRS*		*HRS*		*HRS*			
Program Director - Senior	1880	$218	$409,840		$ -		$ -					$ -	$ 409,840.00
Solutions Architect - Senior	2820	$200	$564,000		$ -		$ -					$ -	$ 564,000.00
Governance Lead - Senior	2820	$195	$549,900		$ -		$ -		$ -		$ -	$ -	$ 549,900.00
Security Architect - Senior	1880	$200	$376,000		$ -		$ -		$ -		$ -	$ -	$ 376,000.00
Scurm Master - Senior	3760	$174	$654,240		$ -		$ -		$ -		$ -	$ -	$ 654,240.00
Training and Adoption Lead - Mid	1880	$150	$282,000		$ -		$ -		$ -		$ -	$ -	$ 282,000.00
System Engineer - Senior	3760	$175	$658,000		$ -		$ -		$ -		$ -	$ -	$ 658,000.00
UI/UX Designer - Senior	3760	$165	$620,400		$ -		$ -		$ -		$ -	$ -	$ 620,400.00
Full-Stack Developer - Senior	3760	$176	$661,760		$ -		$ -		$ -		$ -	$ -	$ 661,760.00
Full-Stack Developer - Mid	3760	$155	$582,800		$ -		$ -		$ -		$ -	$ -	$ 582,800.00
Full-Stack Developer - Junior	3760	$135	$507,600		$ -		$ -		$ -		$ -	$ -	$ 507,600.00
Total Direct Labor			$5,866,540	0	$ -	0	$ -	0	$ -	0	$ -	$ -	$ 4,892,700.00
II. Indirect Costs													
Direct Labor Overhead Cost	Rate (%):	0%	$0										
Direct Labor Fringe Cost	Rate (%):	0%	$0		$ -		$ -		$ -		$ -	$ -	$ -
III. Subcontractors/Consultants/Contracted Staffing													
JDSAT, Inc.	*HRS*	*Rate*		*HRS*		*HRS*		*HRS*		*HRS*			
Chief Data Scientist	2585	$171	$442,035		$ -		$ -		$ -		$ -	$ -	$ 442,035.00
Data Architect- Senior	2585	$208	$537,680		$ -		$ -		$ -		$ -	$ -	$ 537,680.00
Synthetic Data Engineer- Mid	2585	$142	$367,070		$ -		$ -		$ -		$ -	$ -	$ 367,070.00
AI Engineer- Mid	2397	$142	$340,374		$ -		$ -		$ -		$ -	$ -	$ 340,374.00
Data Analyst - Junior	2820	$83	$234,060		$ -		$ -		$ -		$ -	$ -	$ 234,060.00
Software Engineer- Senior	3525	$198	$697,950		$ -		$ -		$ -		$ -	$ -	$ 697,950.00
Software Engineer- Junior	3290	$112	$368,480		$ -		$ -		$ -		$ -	$ -	$ 368,480.00
DevOps Engineer- Mid	1692	$143	$241,956		$ -		$ -		$ -		$ -	$ -	$ 241,956.00
Total Subcontractors/Consultants			$3,229,605	0	$ -	0	$ -	0	$ -	0	$ -	$ -	$ 3,229,605.00
IV. Direct Materials (provide description and costs basis)													
Curriculum Development & Training			$0									$0	$0
			$0									$0	$0
			$0									$0	$0
			$0									$0	$0
			$0									$0	$0
			$0									$0	$0
			$0									$0	$0
Total Direct Materials			$0		$0		$0		$0		$0	$0	$0
V. Travel			$19,890									$0	$19,890
VI. Other Direct Costs *(list type and description such as venue rental costs and cost basis)*													
Cloud Infrastructure Services (compute/storage/network/security)			$364,365									$ -	$ 364,365.00
Development tooling (AI assisted dev tools)			$19,296									$ -	$ 19,296.00
												$ -	$ -
			$0									$ -	$ -
			$0									$ -	$ -
Total Other Direct Costs			$383,661		$ -		$ -		$ -		$ -	$ -	$ 383,661.00
VII. General & Administrative Expense	Rate (%):		$0										
Total PROJECT COSTS			**$9,499,696**		**#REF!**		**#REF!**		**#REF!**		**#REF!**	**#REF!**	**#REF!**